

PAPER • OPEN ACCESS

## Crust & Crumb sampling: efficient sampling design for high dimensional parameter recovery


To cite this article: R J Ward and F Orihuela-Espina 2025 *Meas. Sci. Technol.* **36** 065410

View the [article online](#) for updates and enhancements.

You may also like


- [Multilayer manipulated diffraction in flower beetles \*Torynorrhina flammea\*: intraspecific structural colouration variation](#)  
C X Song, F Liu, Y H Hao et al.
- [Tungsten Disulfide Wrapped  \$WO\_3\$  Based Thin Film Transparent Electrode Material for Electrochromic Device Application](#)  
Rahul Sharma, Nihal, Mamta Sharma et al.
- [Optimisation of Preparative and Performance Parameters on Electrochromic Properties of Electrochemically Deposited Tungsten Oxide Films](#)  
Michael G. Hutchins, Nasser A. Kamel, Nabila El-Kadry et al.

Join the Society  
Led by Scientists,  
for *Scientists Like You!*



The  
Electrochemical  
Society

Advancing solid state &  
electrochemical science & technology



# Crust & Crumb sampling: efficient sampling design for high dimensional parameter recovery

R J Ward\*  and F Orihuela-Espina 

School of Computer Science, University of Birmingham, B15 2TT Birmingham, United Kingdom

E-mail: [rjw207@student.bham.ac.uk](mailto:rjw207@student.bham.ac.uk)

Received 28 February 2025, revised 9 May 2025

Accepted for publication 29 May 2025

Published 9 June 2025



CrossMark

## Abstract

Colouration maps offer a look-up table solution to inverse reconstruction problems in optical imaging. The initial generation of these maps is, however, expensive due to combinatorial explosion or exposed to large errors if the number of forward simulated points is compromised. Here, Crust & Crumb, a pseudo-random gradient-weighted sampling design is put forward aiming to circumvent the combinatorial explosion whilst keeping the reconstruction error at bay. A synthetic exponential model alike the Beer–Lambert law is constructed to establish the validity of the new sampling design, and a finite element diffusion model exemplifies its feasibility against spectroscopic data. A real-data example illustrates haemodynamic reconstruction against a typical linear least-squares approach. When compared to alternative sampling designs, Crust & Crumb affords the same reconstruction error as a formal gradient-based approach, and posit that it exhibits the scalability needed for applications such as broadband near-infrared spectroscopy. The algorithm is agnostic to the domain of application, and hence, it may have implications for other applications of interpolation.

Keywords: broadband near-infrared spectroscopy, sampling, colouration maps

## Acronyms

$\mu_a$	Absorption coefficient	LUT	Look-up table
$\mu'_s$	Reduced scattering coefficient	MBLL	Modified Beer–Lambert law
bNIRS	Broadband near-infrared spectroscopy	NIRS	Near-infrared spectroscopy
C&C	Crust and Crumb	O2Hb	Oxygenated haemoglobin
CCO	Cytochrome-C oxidase	RBFs	Radial basis functions
CM	Colouration map	RTE	Radiation transport equation
FEM	Finite-element modelling	SNR	Signal to noise ratio
FF	Fat fraction	WF	Water fraction
HHb	Deoxygenated haemoglobin		
LL	Log-likelihood		

\* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

## 1. Introduction

In spectroscopy, one aims to recover a set of parameters characterizing the matter based on the fit of observations with a forward model. This inverse reconstruction can be solved with a CM. A CM is a LUT table describing all or a subset of possible colours that may be observed as a consequence of light transport through matter as a function of parameters characterizing the matter. When not all colours are available, the rest are

interpolated. CMs provide fast solutions to inverse problems and can achieve any arbitrary degree of precision, featuring in state-of-the-art systems enabling scalable portable platforms (Stillwell *et al* 2021).

In CMs, the efficiency in the reconstruction is possible because the computational cost of generating the mapping happens one-off and it is paid before usage. The trade-off between reconstruction error and cost of map generation is largely determined by the number of entries on the LUT, i.e. those forwardly simulated. For a fixed number of samples and hence similar cost, their loci over the parameter space also greatly influence the reconstruction error. The decision of where to sample is known as the *sampling design problem* (Olea 1984).

Sampling designs optimal for minimizing reconstruction error are known for certain scenarios (Wei *et al* 2012). However, we are not aware that a universal solution has been established. Also, sacrificing numerical optimality for computational scalability is convenient in many applications. For instance, uniform sampling speeds up interpolation (Zhang *et al* 2021) and opens up the range of available interpolation methods, sometimes at the cost of reconstruction error. The choice of sampling design depends on the required interpolation error and the computational or monetary cost of sampling, e.g. geological or agricultural sampling. In the absence of physical and/or economic restrictions, sampling designs typically traverse grid-like spaces with as few samples as possible (Mersereau 1979). The sampling design is also influenced by the interpolation technique. For instance, in Kriging, a simplicial lattice is optimal for minimising interpolation error under certain assumptions (McBratney *et al* 1981).

In NIRS and more specifically bNIRS, reducing the computational cost of physiological parameter reconstruction is important for the advancement of real-time imaging via diffuse light methods. LUT-based techniques are one such approach, with their implementation being a field of active research (Abdalsalam *et al* 2024, Gómez and Roblyer 2024). Methods to optimise the use of LUTs are also being explored, from avoiding the need of interpolation through resolving two optical properties from a highly dense grid (Angelo *et al* 2016) to multi-stage LUTs (Yang 2021). In all of these the problem of sampling is not the focus.

Recently, popular work has been on the use of deep learning (Yoo *et al* 2020, Eastmond *et al* 2022). Analogous to CM, deep learning models once trained, are fast at producing reconstructions. This front-loading of the computational time can be seen as analogous to the time required for generating forward models to populate a CM. The two are not necessarily mutually exclusive either, with data generation through LUT and Monte Carlo simulations being used for neural-network training (Kao and Sung 2022). Moreover, it has been shown that the amplitude and phase, optical properties, and chromophores of frequency-domain (FD-NIRS) measurements can be acquired, processed, and displayed at speeds of 36 608 Hz, 17 891 Hz, and 10 211 Hz (Stillwell *et al* 2021), significantly faster (1000x) than iterative linear optimisation methods.

Notwithstanding, in contrast with deep models, with CMs the assumptions made are more easily known and interpreted.

## 2. Background theory

A CM  $\Pi \rightarrow (\Lambda \rightarrow \mathbf{I})$ , or  $\Pi \rightarrow \mathbf{I}(\Lambda)$  for convenience, is a mapping encoding radiation transport as perceived at the location of a detector.

where:

- $\Pi = \{\pi_n\}$  is the space of parameters, often in  $\mathbb{R}^n$ .
- $\Lambda = \{\lambda_k\}$  is the space of wavelengths. In nature  $\Lambda$  exists in  $\mathbb{R}^\infty$ , but discretized approximations represents  $\Lambda$  in  $\mathbb{R}^k$ . In bNIRS,  $k \geq n$ .
- $\mathbf{I}(\Lambda) = \{I(\lambda_k)\}$  is the space of intensities at the different wavelengths. The dimensionality of  $\mathbf{I}$  matches that of  $\Lambda$ , i.e.  $\mathbb{R}^\infty \approx \mathbb{R}^k$ .

In other words, a combination of intensity measurements at particular wavelengths corresponds to a particular set of physiological parameters, as seen in figure 1.

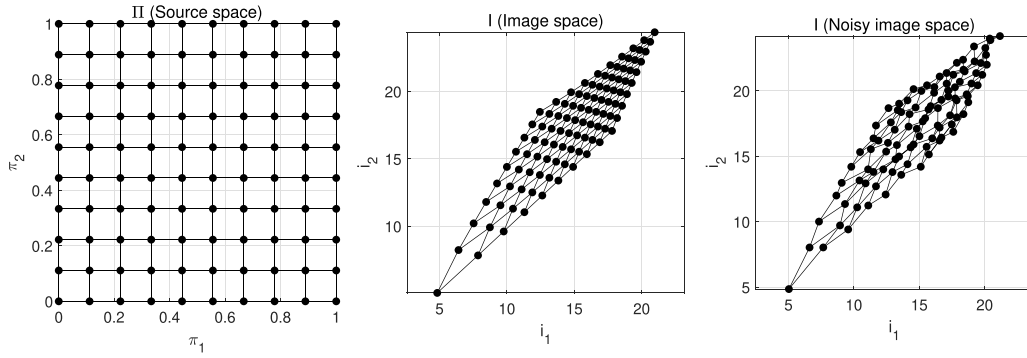
A CM is ad-hoc for a given geometrical medium in which light propagates. This idealized geometrical medium is referred to as the *geometrical model*. In bNIRS, this corresponds to some biological tissue. In such an application, the parameters serve to instantiate the relative concentrations of chromophores but other quantities, e.g. thickness of a given tissue layer or other optical properties, may also be treated as parameters.

We lack a general analytical solution to the RTE for an arbitrary complex geometry. There exists a number of approximations to the RTE (Arridge and Hebden 1997), referred to as the *optical model*. A common choice in NIRS applications is the MBLL (Delpy *et al* 1988). To obtain specific solutions, a range of numerical solvers referred to as the *numerical model* are available e.g. ad-hoc analytical (Fantini *et al* 1994), FEM (Dehghani *et al* 2009), Monte Carlo (Wang *et al* 1995), etc. Together, the geometrical model, the optical model, and the numerical model form the *forward model* and the CM solves the inverse problem of image reconstruction.

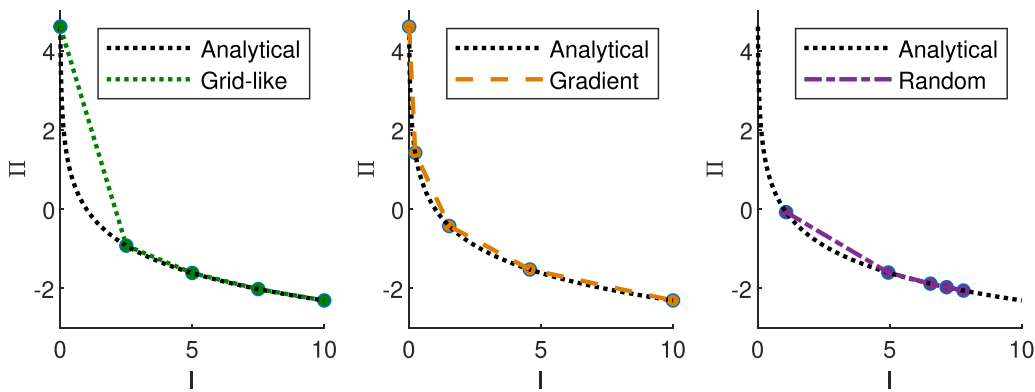
### 2.1. Inverse model considerations

The CM is simply a collection of parameter vectors  $\pi$  forwardly simulated. Given an spectroscopic measurement  $I(\Lambda)$  or colour readily available in the CM, the generating parameter vector  $\pi$  is simply found by looking up the CM table.

For colours not in the LUT, interpolation is used. Regularities (or absence) in the sampling design enables (or disables) the applicability of certain interpolation methods. Given a set of observations at random locations, sometimes referred to as unstructured grid, methods such as linear, RBFs, and natural neighbour interpolation are all amenable. They differ in the rationale and assumptions of the interpolant, but any suit our purposes here. Such interpolation methods will



**Figure 1.** A fictitious non-linear mapping from parameter space  $\Pi$  (with individual parameters  $\pi$ ) to some image space  $I$  for two measurement domains or wavelengths  $i$ . Image inspired by Herrera-Vega *et al* (2017).



**Figure 2.** The negative logarithm for a parameter  $\Pi$  can be estimated through different sampling methods across observations  $I$ . Both grid-like and gradient based methods cover the space, whilst random sampling may not.

typically take the form of an estimation formula in equation (1) (Li and Heap 2014)

$$\hat{\pi}(i_0) = \sum_{l=0}^M \lambda_l \pi(i_l) \quad (1)$$

in which we take some observed value  $\pi(i_l)$  at a given location accompanied with a weighting  $\lambda_l$ , and estimate  $\pi$  at our point of interest  $i_0$  for  $M$  sampled points used in the estimation.

We expect RBF interpolation to outperform the other two methods mentioned for the case of higher dimensionality interpolation. Interpolation methods relying on triangulation of the space or nearest-neighbours implicitly suffer as the dimensionality of the space increases significantly, often becoming either ill-defined or computationally intractable. RBF interpolation is both more suited to managing the curse of dimensionality, and algorithmically simpler for application over large datasets (Ávila-Sansores *et al* 2020).

To guarantee well-posedness and, hence, uniqueness in the reconstruction, the CM ought to be injective, which can be afforded by guaranteeing monotonicity.

## 2.2. Sampling design

Regardless of the forward model, algorithms to generate a CM differ in the sampling design. Common sampling designs can be coarsely summarised into;

- **Grid-like:** Lattice-based sampling of the parameter space. The grid does not need to be equi-space nor simplicial based (Hemingway 2002), but for simplicity, we shall stick to equi-space squared grids here.
- **Random or distribution based:** Random sampling across the parameter space, often uniformly across the parameter space, but where covariance of the parameter space is available, other distributions have been considered.
- **Gradient-based:** A variant of the grid-like for cases when the explicit form of the forward model is known. The sampling design is denser where the gradient of the forward model with respect to it is parameters is smaller to reduce the reconstruction error, and vice-versa, the regions in which the gradient is larger are more sparsely covered.

The effects of each sampling regime are naively illustrated in figure 2.

Both grid-like and gradient-based sampling schemes are fundamentally tied to the dimensionality of the space  $d$ , where given some number of points chosen per axis  $n$  the total sampled points generated is  $n^d$ . We refer to this as the combinatorial explosion. This is in opposition to random sampling designs, where a number of points can be determined by the user and their available computational resources.

For readers interested in the mathematics of designing optimal sampling strategies we suggest (Bellhouse 1984). For

those interested in spatial interpolation techniques we recommend (Li and Heap 2014).

### 3. Methods

When using the MBLL in spectroscopy, describing an exponential decay of light measured given increasing absorption, in principle only 3 points are needed to reconstruct the absorption curve. In practice, a larger number of sampling points are often used in CMs for robustness against numerical inaccuracies and noise. Even with the minimal number of points across each dimension the combinatorial explosion quickly dominates the scene. Grid-like and gradient-based approaches suffer from this combinatorial explosion.

As mentioned above, given a sampling for which there are  $n$  points chosen per parameter dimension  $d$ , the number of total sampled points will be  $n^d$ . If one aims to model O2Hb, HHb, WF, FF, CCO, melanin content (contributing to the Absorption Coefficient ( $\mu_a$ )) and the Reduced Scattering Coefficient ( $\mu_s'$ ) even with only 5 samples per parameter for  $k$  wavelengths there will be  $5^{n=7} \times k = 78125 \times k$  simulations needed to form the (possibly still sparse) CM when using a combinatorial approach. For instance, for the case of 200+ wavelengths common in bNIRS, and an estimated simulation time of 1.6 s per simulation, this would take >300 days to generate. Further, this assumes a homogeneous tissue. It can be appreciated how multilayered models soon reach prohibitive numbers of sample points (and simulation time).

Random sampling is less affected by the combinatorial explosion, even though not immune to the curse of dimensionality, with the number of sampled points chosen per the users requirements and resources. However, the random approach does not guarantee covering the boundaries of the parameter space, and hence inconvenient extrapolation may be needed. The price to pay is often a larger reconstruction error near those boundaries, and less rigorous assumptions for the edge of the space.

#### 3.1. Crust & Crumb

Ideally, we want the coverage of a grid-like sampling, the reconstruction error of gradient-based sampling, and the computational flexibility of random sampling. With such aim, we propose C&C as an approach for sampling a given space. C&C exploits the gradient weighting whilst avoiding the combinatorial explosion inherent in grid-like and gradient-based methods by allowing the user to determine the density with which the inner space is sampled. The aim is achieved by identifying the boundaries of the space (the Crust) and pseudo-randomly sampling the inner region (the Crumb) with respect to the gradient of the space.

The Crust is established using a two-step process. First, the straight lines between the interval limits of each parameter's domain are sampled (*hull*). These interval limits can be chosen to mimic physiologically expected ranges for the case of NIRS. This can be sampled in different manners: grid-like, gradient-weighted, random, or just the vertices. Then

---

#### Algorithm 1. Crust & Crumb sampling design.

---

**Require:**  $lims \in \mathbb{R}^{n \times 2}$ : Ordered parameter boundaries [min, max]  
**Require:**  $crustPoints \in \mathbb{N}$ : Number of points in the crust sampling design  
**Require:**  $crumbPoints \in \mathbb{N}$ : Number of points in the crumb sampling design  
**Require:**  $bubbleSize \in [0, 1] \in \mathbb{R}$ : Normalized bubble size  
**Require:**  $rate \in \mathbb{R}$ : Gradient steepness control  
**Require:**  $\alpha \in \mathbb{R}$ : Convexity/Concavity of  $\hat{\Pi}$   
**Ensure:**  $\hat{\Pi}$ : Sampling design, set of samples to be acquired

- 1: //Generate crust
- 2:  $boundaryPoints \leftarrow hull(lims, crustPoints)$
- 3:  $crust \leftarrow rehull(boundaryPoints, \alpha)$
- 4:  $\hat{\Pi} \leftarrow crust$
- 5: //Generate crumb
- 6:  $totalPoints \leftarrow |\hat{\Pi}| + crumbPoints$
- 7:  $minAllowableDistance \leftarrow (lims(2) - lims(1)) \times bubbleSize$
- 8: **while**  $|\hat{\Pi}| < totalPoints$  **do**
- 9:  $\pi \leftarrow lims(1) + (lims(2) - lims(1)) \times rand(1, n)^{rate}$
- 10: **if**  $inPolygon(\pi, crust)$  **then**
- 11: **for all**  $\hat{\pi} \in \hat{\Pi}$  **do**
- 12:  $distances[\hat{\pi}] \leftarrow \|\pi - \hat{\pi}\|_p$
- 13: **if**  $\min(distances) \geq minAllowableDistance$  **then**
- 14:  $\hat{\Pi} \leftarrow \hat{\Pi} \cup \{\pi\}$
- 15: **end if**
- 16: **end for**
- 17: **end if**
- 18: **end while**

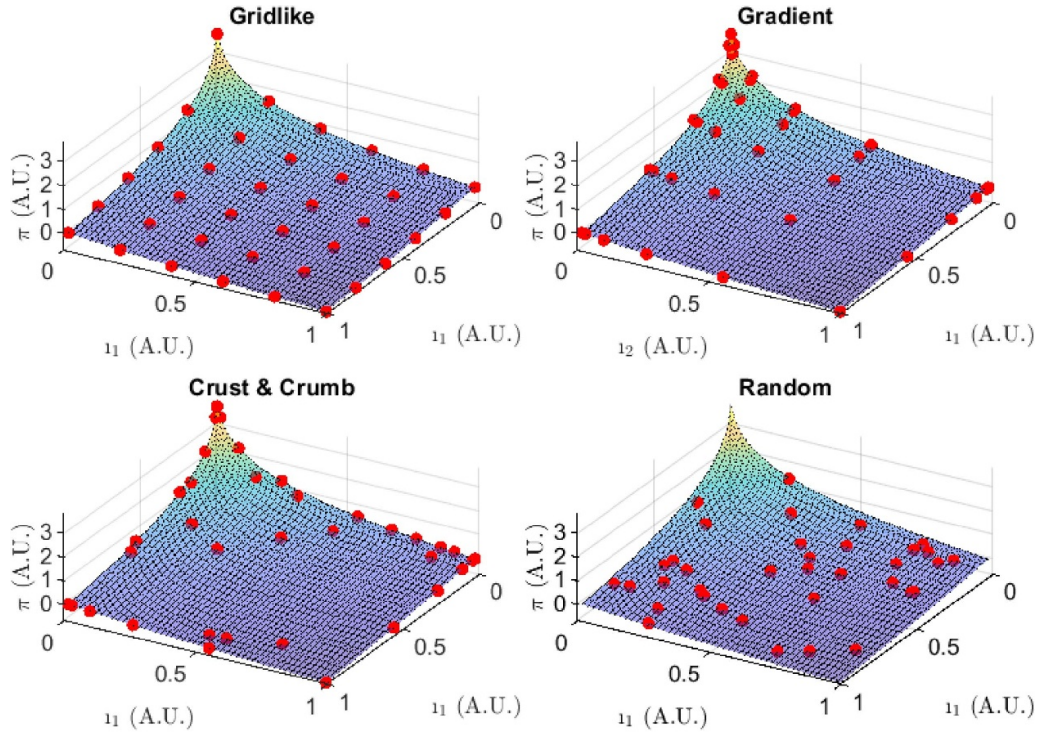
---

some hulling algorithm, e.g. the convex or concave hull (via a method such as alpha-shapes), is forced to (inversely) relocate according to the desired convexity or concavity of the user (*rehull*). We implement this by allowing the hull to become slightly concave, forcing points along the boundary edge to be identified as the crust.

The inner subspace (Crumb) is sampled by weighting a random sample inversely proportional to the absolute gradient of the forward model because any constant perturbation is more impactful as the absolute gradient decreases. We therefore need a greater number of sampled points within that region to ensure interpolation error is low. The tuning of the gradient function depends on the forward model. For a given exponential distribution the bounds and rate of the function steepen or lessen the descent to match the sampling to the model.

Finally, C&C reduces the redundancy of random forward modelling by introducing a minimum spacing requirement, a 'bubble', in which no other point is sampled, and reduces redundancy of forward modelling points occupying very similar spaces. The choice of bubble size is arbitrary to the algorithm and dependant on the space itself. The C&C algorithm is given in algorithm 1. For intuition, C&C is visually compared to other sampling designs in figure 3.

Throughout this paper,  $bubbleSize$  was set to  $0.05 \times$  observation value range,  $hull$  was operationalized using a gradient-based approach, and  $\alpha$  was set to generate the convex hull for *rehull*. The  $rate$  was set to Napier's constant  $e$ . The total number of forward modeled sampling points was set to match



**Figure 3.** Different sampling approaches for an exponential decay model, recovering a parameter value  $\pi$  given two measurements  $i$ . Units are arbitrary.

the combinatorial points generated by equidistant and gradient sampling for testing to facilitate comparison of reconstruction error. Distances were calculated in the Euclidean space, i.e.  $p = 2$ , but for other Lebesgue spaces, one can use a different  $p$ . The function *inPolygon* (Alg line 10) can be implemented using Delaunay's triangulation and is readily available in many programming languages.

A rationale for C&C escaping the combinatorial explosion can be found in [appendix](#).

### 3.2. Synthetic model

A synthetic model was created for validation purposes. Our forward model (proxy of the MBL) is an exponential decay  $\mathbf{I} = e^{(-\Pi)}$ , in which increasing chromophore concentrations  $\pi$  lead to exponential decreases in remitted light  $I$ . This represents the instance that the larger the concentration of a chromophore the lower the intensity of light (at any specific wavelength). The parameter value was given by  $\pi = \sum_k -\ln(I_k)$ .

To understand the weighting of the sampling we look at the reconstruction error, a quantity proportional to the interpolation error. The reconstruction error will be inversely proportional to the gradient of the forward model function  $\pi \rightarrow I(\lambda_i)$ . As the gradient of this function decreases the potential reconstruction error grows, and vice versa. The inverse of our forward model, the negative logarithm, describes the behaviour of our observed values with respect to our interpolant, and aids the intuition behind the direction of the weighting of the sampling.

Interpolation was only applied on a single direction, hence the summation of observation to a single parameter dimension. We observe multivariate interpolation in section 4.7.

## 4. Synthetic results

All work was carried out in Matlab R2024a. The experiments were run on a 13th Gen Intel(R) Core i7-13 700 H Processor, with a NVIDIA GeForce RTX 4060 Laptop GPU and 16.0GB of RAM. All code is available on Github<sup>2</sup>.

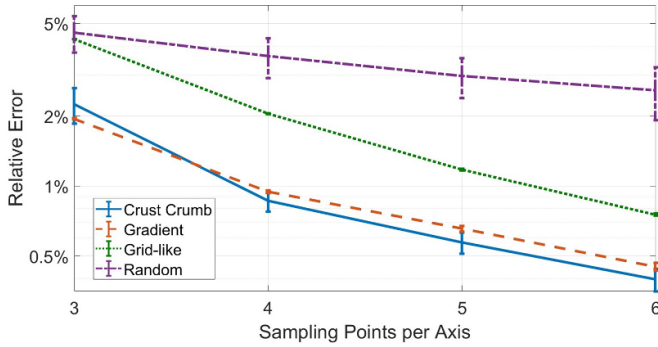
When enabling fair comparisons we look to match the number of sampled points across the different sampling regimes. One principle benefit of C&C sampling is that given a generated Crust, the inner points can be generated on an ad-hoc basis, with any number of points bounded by computational cost. This benefit is ignored for the sake of fair comparison across methods, but if C&C can match the reconstruction error of other designs it should be clear the benefits of sampling beyond methods limited to dimension-dependant regimes.

In the subsequent experiments, interpolated  $\pi$  were compared across a uniformly sampled set spanning  $\Pi$  of 200 test points, for 50 iterations. Simulations were done in Matlab. We report relative error. The region of interest is  $1.0 \leq I \leq 100.0$ , and a perturbation of fixed amplitude 0.1 and random direction was applied to each observation  $I$ . Since random sampling does not cover the boundaries, to facilitate comparison, we pick some nominal boundaries from the synthetic model generation.

<sup>2</sup> <https://github.com/ward-compsci/Crumb-and-Crust>.

**Table 1.** Effect of interpolation method on reconstruction error.

Interpolation	$\beta$	SE	95% CI	$t$ -statistic	$p$ -Value
Intercept (Linear)	0.013	0.0041	[0.0055 0.021]	3.3	$9.7 \times 10^{-4}$
Change (RBF)	$-1.1 \times 10^{-4}$	$1.3 \times 10^{-4}$	$[-3.7 \times 10^{-4} 1.5 \cdot 10^{-4}]$	-0.81	0.42
Change (Natural)	0.0029	$1.3 \times 10^{-4}$	[0.0027 0.0032]	22	$< 1 \times 10^{-10}$

**Figure 4.** Reconstruction error for a 2D observation set using RBF interpolation, with respect to points sampled per axis. A log-scale is used for the y-axis.

#### 4.1. Choice of the interpolation method

To account for the effect of the interpolation method, we compared the performance across 3 interpolation methods as we increased the number of sampled points per dimension. 50 replications for 3, 4, 5 and 6 sampled points per dimension were run as described in section 4. A linear mixed-effects model was fitted with maximum likelihood estimation, with random hierarchical intercepts for dimension, sampling and lattice ( $N_{\text{obs}} = 4,800$ ,  $LL = 19,837$ ). RBF and linear interpolation performed statistically similar whilst natural interpolation performed more poorly, as shown in table 1. As mentioned earlier, from existing literature (Ávila-Sansores *et al* 2020) we expect RBF interpolation to be more suitable in the general case of high dimensional interpolation. Following this we will continue to use RBF interpolation in further experiments.

#### 4.2. Choice of the sampling design

We run 50 replications with perturbations for RBF interpolation.  $n = 2, 3, 4$  measurement dimensions are used to further account for scalability. Similar to section 4.1, a mixed-effects model is fit to the reconstruction error and the effect of sampling designs are observed accounting for points sampled and dimensionality of the dataset ( $N_{\text{obs}} = 2,400$ ,  $LL = 9756.6$ ).

The reconstruction error with respect to points sampled for a 2D observation set is illustrated in figure 4. Random sampling and C&C have 2 sources of variability; the perturbation and their inherent random nature, naturally leading to larger dispersion than the deterministic strategies (grid-like and gradient) for which the variability is only due to the perturbation. This is reflected in the error bars of figure 4,

and quantified in table 2. The designs incorporating random elements have a standard deviation of typically an order higher than deterministic designs. Additionally, the variance of C&C sampling trends towards the deterministic methods as the points sampled increases, justifying the use of a redundancy check to force effective sampling of the available space.

As anticipated, C&C and gradient-based sampling outperformed the other two designs (table 3). A contrast test showed no statistical difference between the gradient and C&C sampling designs,  $F(1, 2396) = 0.009, p = 0.92$ , suggesting our proposed sampling offers the same reconstruction capabilities as the formal gradient-based design, without being restricted to a design susceptible to combinatorial explosion.

#### 4.3. Choice of bubble size

As mentioned in section 3.1 the choice of bubble size to reduce redundancy is dependant on the topology of the space itself, alongside the dimensionality and number of points sampled. The effect of increasing the bubble (where the redundancy check is the bubble  $\times$  observation value range) for a 2 dimensional synthetic dataset is provided in figure 5, for 50 repetitions. For the case of few sampled points this does not significantly affect the reconstruction error, but as the number of sampled points increases maximising the redundancy space improves the reconstruction error. There is of course an upper-limit for the size of the bubble; the size needs to be small enough to allow all the points to fit within the bounded space. Compared to other parameters of sampling design the choice of bubble size has a relatively small effect on reconstruction error for these particular cases.

#### 4.4. Effect of noise

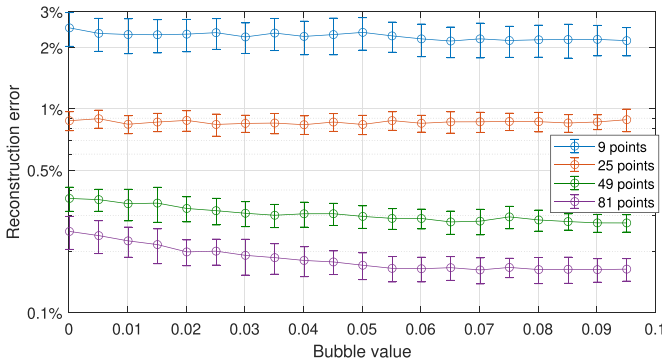
As the noise of a system increases, the reconstruction error will worsen. Understanding the response of the sampling strategy in light of different noise scenarios allows discussing the effectiveness of a design. The behaviour of each sampling design with respect to an increasing perturbation over 50 repetitions is illustrated in figure 6. For the case of 3 points per axis, grid-like and random sampling both exhibit the same behaviour, with gradient and C&C sampling consistently offering lower reconstruction error. This is in keeping with initial theory, where sampling regions more affected by noise results in a better reconstruction. As the number of sampled points increases the behaviour of all-bar random trends to the same model, indicating the space and simplistic model is sufficiently sampled. The randomly sampled CM systematically performs worse than the other designs.

**Table 2.** Standard deviation of reconstruction error for a 2D dataset, given sampling design and points per axis.

Points per axis				
Sampling	3	4	5	6
Grid-like	$4.3 \pm 0.0066\%$	$2.0 \pm 0.0089\%$	$1.1 \pm 0.0092\%$	$0.75 \pm 0.0085\%$
Gradient	$1.9 \pm 0.011\%$	$0.94 \pm 0.014\%$	$0.66 \pm 0.018\%$	$0.45 \pm 0.018\%$
Random	$4.6 \pm 0.82\%$	$3.6 \pm 0.70\%$	$3.0 \pm 0.58\%$	$2.6 \pm 0.66\%$
C&C	$2.2 \pm 0.39\%$	$0.86 \pm 0.087\%$	$0.57 \pm 0.060\%$	$0.40 \pm 0.044\%$

**Table 3.** Effect of the sampling design on reconstruction error, accounting for dimension and number of points.

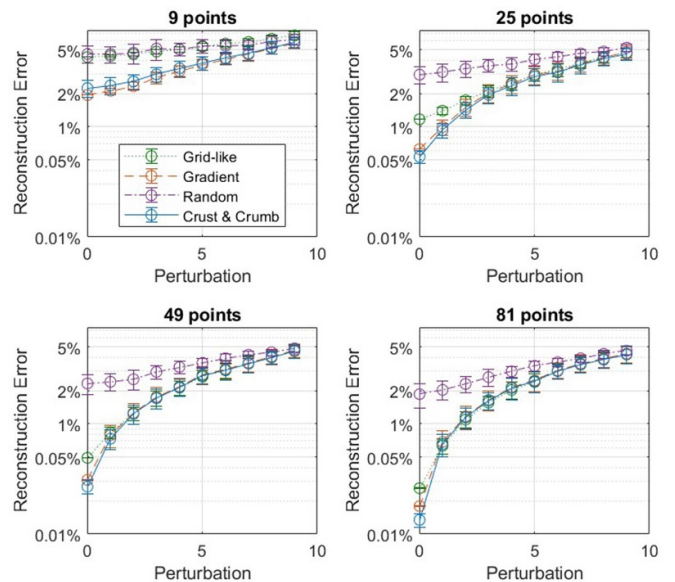
Sampling	$\beta$	SE	95% CI	$t$ -statistic	$p$ -Value
Intercept (Grid-like)	0.0092	0.0036	[0.0022 0.016]	2.6	0.0095
Change (Gradient)	-0.0036	$2.4 \times 10^{-4}$	[-0.0041 - 0.0032]	-15	$< 1 \times 10^{-10}$
Change (Random)	0.011	$2.4 \times 10^{-4}$	[0.011 0.012]	48	$< 1 \times 10^{-10}$
Change (C&C)	-0.0036	$2.4 \times 10^{-4}$	[-0.0041 - 0.0031]	-15	$< 1 \times 10^{-10}$



**Figure 5.** The effect of bubble size on relative interpolation error using RBF interpolation across a 2D C&C sampled dataset. A log-scale is used for the y-axis.

4.5. Effect of Crumb density

With the current implementation of our algorithm the crust is determined by hulling a gradient-based mesh and then populating the inner space with Crumbs. We have already shown in section 4.2 that C&C sampling can match the reconstruction error of gradient-based sampling and exceed that of random and grid-like sampling. As such we may now observe the behaviour of the reconstruction error given different starting Crusts and an increasing number of Crumbs. Four simulations were run for 50 cycles each, increasing the total number of points used for the sampling, with the results shown in figure 7. Increasing the number of Crumbs provides a large initial improvement on the reconstruction error but this reduces as the space is filled out. Of note is that for  $n$  points in  $k$  dimensions, we have the flexibility to sample a number of points between  $n^k$  and  $(n + 1)^k$  points, a limitation of gradient and grid-like sampling. As such we can make use of the additional error reduction illustrated given computational power. This approach is inline with our theory in appendix.

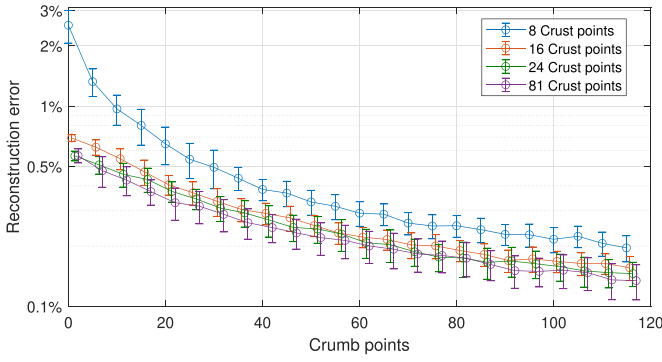


**Figure 6.** Effect of increasing perturbation on 2D synthetic dataset using RBF interpolation given increasing points sampled per axis. The perturbation is an absolute value where the test range is  $1.0 \leq I \leq 100.0$ . A log-scale is used for the y-axis.

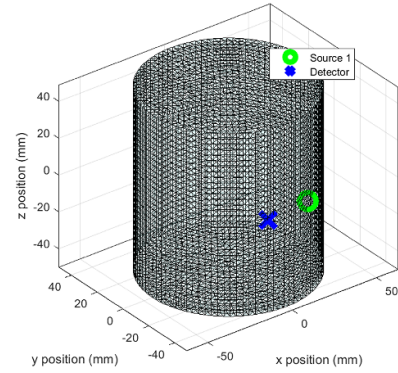
4.6. Effect of undersampling the colouration map

As mentioned in section 4, fair comparison is only possible if one uses the same number of points. In principle, in order to observe differences, one could eliminate  $k$  points from both approaches, and expect that the ‘gaps’ in the sampling suffered by the gradient grid due to its rigid loci would harm the reconstruction. Whilst this is a fair comparison with respect to the points used, we do note that the assumptions of maintaining boundaries for grid-like, gradient and C&C sampling is not upheld.

Figure 8 explores the effect of randomly undersampling each sampling design for a fairly densely sampled CM.



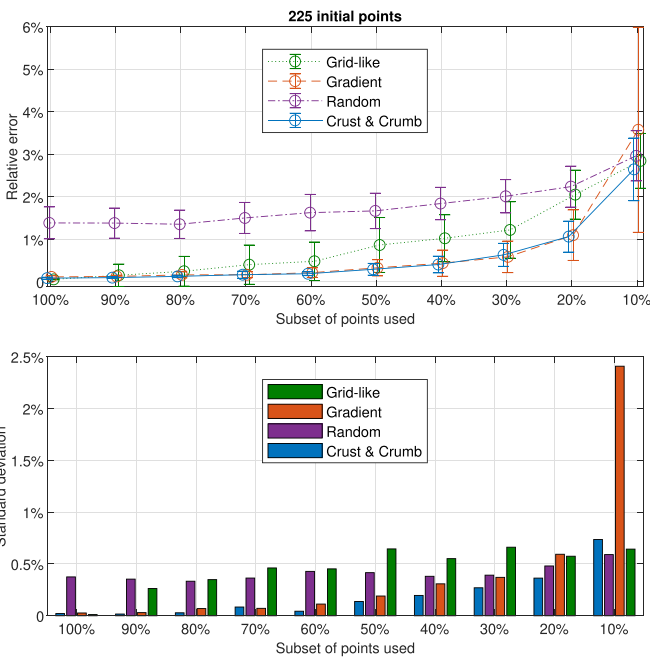
**Figure 7.** Effect of increasing the number of Crumb (or inner) points on the relative reconstruction error, for varying numbers of points used to define the Crust (or boundary). A log-scale is used for the y-axis. A jitter is applied along the x-axis for readability.



**Figure 9.** Forearm mesh used in simulations, with source and detector locations matching those of the inhouse system.

**Table 4.** Tissue property bounds used for the CM map generation.

Chromophore	Lower bound	Upper bound
O2Hb	20 $\mu\text{M}$	80 $\mu\text{M}$
HHB	10 $\mu\text{M}$	70 $\mu\text{M}$
WF	0.1%	0.7%
FF	0.1%	0.7%
Scattering parameter	Lower bound	Upper bound
$\mu'_s$	0.4 $\text{mm}^{-1}$	1.5 $\text{mm}^{-1}$



**Figure 8.** Top: Reconstruction with a randomly sampled subset of a  $15^2$  point CM, over 50 cycles of 200 test points shown. A jitter is applied along the x-axis for readability. Bottom: Dispersion of the above reconstruction. The lower dispersion of C&C suggests that it is more stable than alternatives.

The lower standard deviation (and thus variance) over 50 replications of the C&C sampled design remains consistently below the alternatives. In other words, C&C is more stable to alterations of the loci of the sampled points and less sensitive to undersampling.

#### 4.7. Synthetic forearm measurements

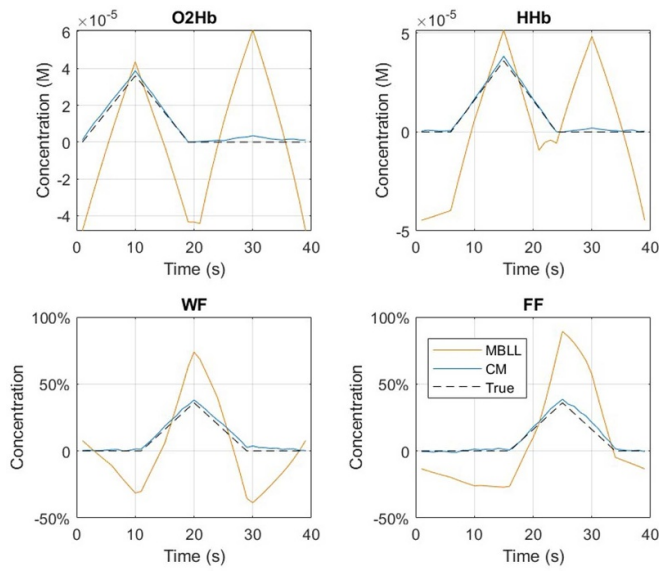
To test the CM reconstruction on a more elaborate synthetic model an FEM theory approach was followed. A 3D cylindrical mesh was generated simulating the forearm in NIRFASTer (Deghani *et al* 2009) for muscle oxymetry, shown in figure 9. Source and detector locations were set to match that of the inhouse broadband system (Ward *et al* 2023).

Physiologically plausible parameter limits were chosen for O2Hb, HHB, WF, FF (the primary absorbing contributors of  $\mu_a$ ), and  $\mu'_s$  from previous literature (Jacques 2013), given in table 4.

While it is of interest to understand metabolism, CCO was not included as it has a minimal effect on absorption compared to the main absorbing chromophores, and hence alleviates the experimental simulation cost. A CM sampled via C&C was generated, and the forward model ran for 211 wavelengths. The resulting size of the CM is  $3992 \times 211$ .

Typical bNIRS measurements are taken as time series measurements, over some tissue surface of interest. To simulate the effect of changing absorbers within the tissue, each contributing chromophore and  $\mu'_s$  was linearly increased and decreased over a course of 20 s per component. Each change started 5 s after the previous to ensure an overlap of varying contributors.

The MBLL (or variations on) is the most widely used reconstruction method within bNIRS, and as such we use it to validate our reconstructions against existing literature. The results of the reconstructions, visually compared against reconstructions via the MBLL, are shown in figure 10. Whilst the absolute values generated from the MBLL are not concerning (a scaling factor such as the differential pathlength factor can resolve these) the CM is shown to be more robust to multiple crosstalk between chromophores. The scattering change was not reconstructed for the figure as the fundamental MBLL does not offer a means to solve for scattering without further terms, yet the impact of a changing scattering factor can be



**Figure 10.** Reconstruction of four chromophores via MBLL and our CM using RBF interpolation, as the ground truth of each is varied over across a time series (O2Hb, HHb, WF, FF).

seen as the MBLL hallucinates oxygenation changes over the final 20 s of the experiment.

### 5. Real measurements


The CM generated in section 4.7 was used for reconstruction of *in-vivo* measurements of the forearm during a cuff occlusion as detailed in section 5.2.

#### 5.1. Data pre-processing

As the CM was designed to measure and reconstruct absolute changes of chromophores, a calibration was carried out to map measurements to the interpolating surface. Measurements were collected from two phantoms with known  $\{\mu_a, \mu_s'\}$  at set wavelengths, and these were used to calibrate the CM. The calibrated data was denoised through a subtraction of a dark-noise measurement taken before the experiment, and smoothed via a 3rd order Savitzky–Golay filter. This removed hot pixels (fixed pattern) and dark noise from our results, and Poisson noise was not deemed an issue due to the large SNR used. To resolve limited phantom calibration availability and subsequent conditioning of the measurement matrices only a subset of wavelengths where interpolated over for these results - 8 wavelengths from 690 nm to 900 nm in 30 nm intervals for O2Hb, and in 20 nm intervals for HHb. This allows us to still work in a high dimensional dataset from the perspective of NIRS analysis and is in keeping with current literature on the selection and number of wavelengths for bNIRS analysis (Arifler *et al* 2015).

#### 5.2. Data collection

An arm cuff occlusion was carried out on an adult male participant, with ethics approved through the University of Birmingham (ERN\_0563). This was measured with the



**Figure 11.** Reconstructed haemodynamics on the adult forearm during a cuff occlusion. Reconstructed with the MBLL (top) and with a C&C sampled CM using RBF interpolation (bottom).

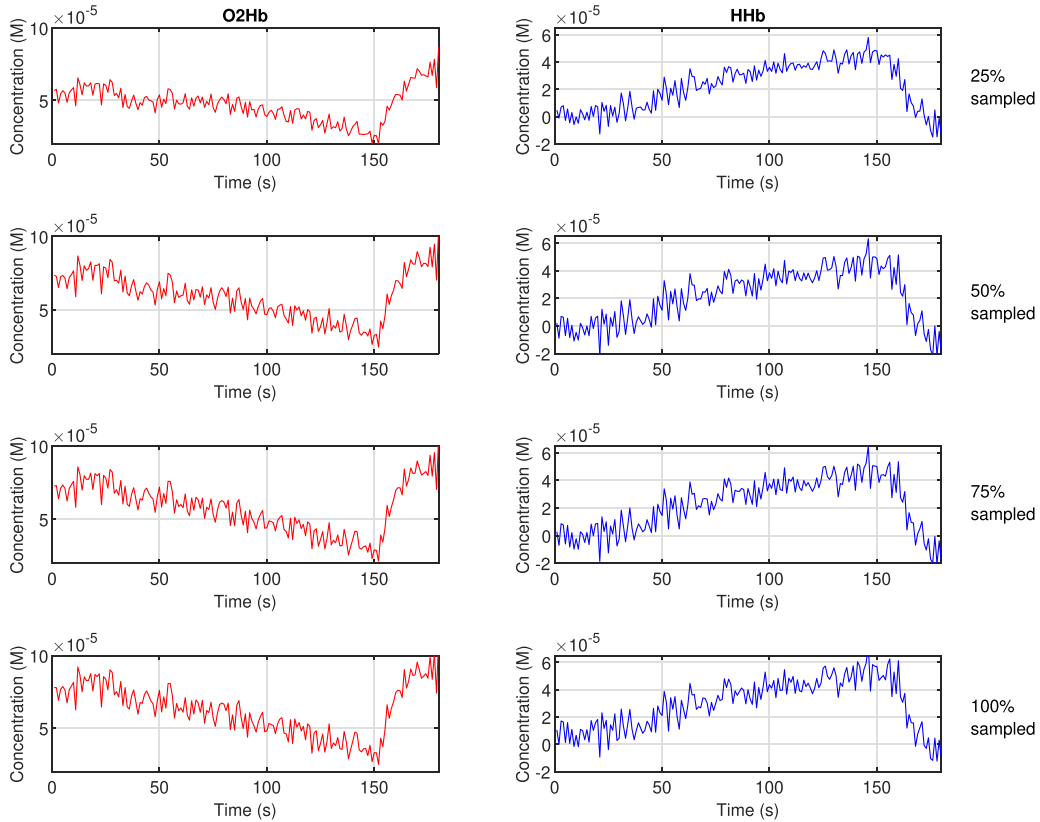
inhouse bNIRS system for a single channel, on the brachioradialis of the left forearm. Cuff occlusion took place over two minutes, with 30 s of resting state data measured either side of the occlusion. An integration time of 1 s was used, which for a single channel gives an effective sampling rate of 1 Hz.

### 5.3. Results

Figure 11 shows the reconstructions of the haemodynamics within the forearm during cuff occlusion. It should be noted that the MBLL is reconstructing relative changes in concentrations, whilst the CM reconstructs the absolute values of haemoglobin concentrations. Hence the MBLL will use a baseline of 0  $\mu\text{M}$ . There is no ground-truth to these results—the changes of oxygenation are not known for these instances, only estimated through changes of light absorption.

As expected, both reconstruction methods capture the dropping oxidation of the haemoglobin as the arterial supply is cut, and the tissue continues to metabolise. The CM predicts analogous changes in O2Hb and HHb, which is physiologically believable—the existing haemoglobin donates its carried oxygen, reducing the absorption of light due to this molecular structure. Meanwhile, the MBLL appears to overestimate the quantity of HHb, estimating a change of double that of the O2Hb.

Non-physiologically sound jumps in both O2Hb and HHb appear in the MBLL reconstruction at the start of occlusion, and the CM reconstruction resolves these. As mentioned earlier in section 5.1, a subset of available wavelengths was used for interpolation and reconstruction. This is to aid in the measurement matrix conditioning, reducing numerical degeneracies. The apparent smoothness of the MBLL results can be compared to those seen in section 4.7, in which the MBLL over-estimates (and thus over-smooths) the time-series.



**Figure 12.** Effect of randomly undersampling the CM on haemodynamic reconstructions for *in-vivo* forearm cuff occlusion measurements, from 25% of the CM (top) to 100% (bot).

The stability of C&C sampling can be empirically observed by taking a random subsample of the generated CM for reconstruction, shown in figure 12. A subsampled of 25% to 100% of the CM is selected for reconstruction. All subsets capture the trend of decreasing blood oxygenation, with the primary disagreements being in the quantity of O2Hb for the severely undersampled case. This is in keeping with the results of section 4.6.

## 6. Discussion

For all practical purposes, it is safe to assume some arbitrarily large yet constant time per forward simulation. Hence, regardless of the sampling design or chosen interpolant, CMs incorporating a larger number of points take proportionally longer to generate. Furthermore, the overhead of the sampling strategy itself is negligible compared to the forward simulation time. Therefore, as long as the number of points of each CM is the same, the generation time will also be approximately the same.

To achieve a certain error threshold, the sampling design that faster converges will necessitate fewer points to be forward simulated to reach such threshold, and hence it will be computationally more efficient. The results above from both a naive exponential model and broadband-like FEM-based diffusion model support the hypothesis that when aiming for

comparable errors, C&C achieves a statistically similar result to that of a regime strictly following the gradient, and thus more efficiently samples the space than grid-like or random sampling, upholding convergent validity. When considered alongside the assumptions of higher dimensionality sampling, C&C offers a computationally efficient alternative for both known and estimated forward models. We can sample the space ad-hoc given computational power without being restricted to the combinatorial sampling constraints of a grid-like design.

As mentioned, C&C is more efficient than strict gradient-based sampling as it can more stably and less sensitively escape the combinatorial explosion to which gradient-based sampling is fated, as illustrated in section 4.6. In the end the error is limited (and bounded) by some unknown optimal sampling, e.g. dictated through some oracle algorithm. The closer one gets to this limit the less differences in effect size would be as this is a capped outcome. This is in combination with greater flexibility—C&C sampling is not restricted to fixed powers of sampling counts. This enables more plastic exploration of the space which can be particularly useful in other domains where the error function itself is not monotonic, noting that the colouring function should be monotonic even if the error distribution itself is not.

The need for a scalable sampling design feeds from computationally cost-effective solutions. C&C sampling makes no assumptions on the dimensionality of the observation or

parameter sets, and as such we predict C&C will succeed in maintaining its effectiveness when taken to higher dimensional datasets. This is crucial for bNIRS applications, in which a large number of physiological parameters are reconstructed from potentially  $> 100$  wavelength intensity measurements. CMs offer a convenient and efficient means of retrieving matter properties. However, one constraint arises from the reliance on a discrete precomputed point cloud, which may not fully capture the optical variations. This limitation becomes more pronounced if the assumption of monotonicity does not hold. Algorithms to correct for deviations of monotonicity exist however (Orihuela-Espina 2005).

C&C offers a strong implementation benefit—with the Crust established, the sampling and forward modelling of points within the Crumb can be completed over multiple visits of the algorithm. The inner space will continue to be filled per the gradient weighting and the redundancy check per point allows computational time to be used as and when it is available. This is of interest when shared servers are in demand (e.g. in University environments).

As new parameters are included, their impact on the transmission of light must be understood to successfully exploit a gradient weighted sampling methodologies. For parameter dimensions in which the forward model is not well-understood uniform sampling still offers an alternative, whilst maintaining C&C sampling along other dimensions.

Future approaches will be to explore the use of sampling designs in explicit bNIRS reconstruction in more depth, understanding how hardware and software can be exploited to assist in reducing the computational burden, whilst incorporating sophisticated interpolation strategies for the inverse problem estimation.

## 7. Conclusions

Motivated by the need to escape the combinatorial explosion when generating a CM for bNIRS, we have proposed and validated a new sampling design, namely Crust & Crumb. C&C offers computational affordability and scalability, whilst maintaining (or even improving on) interpolation error levels compared to other traditional designs. We have shown the generation of CMs for specific tissue (the forearm) and have compared to the MBLL for haemodynamic reconstruction in an arm cuff occlusion. Algorithms such as C&C assist the communities' demand for compact imaging platforms -as reconstruction becomes inherently cheaper-, and beyond bNIRS to domains in which sampling is required over suitable forward models.

## Data availability statement

The data cannot be made publicly available upon publication because they are not available in a format that is sufficiently accessible or reusable by other researchers. The data that support the findings of this study are available upon reasonable request from the authors.

## Acknowledgment

The author would like to thank the University of Birmingham for the scholarship funding to carry out the PhD.

## Appendix. Rational for escaping the combinatorial explosion

The combinatorial explosion represents the upper bound in terms of the number of points forming a colouration map, and hence the basis for comparison and the best case for the potential error. For  $k$  dimensions and  $n$  points per dimension, the combinatorial methods shall require  $M = n^k$  points to be sampled. Given some oracle algorithm to determine the optimal location of these  $M$  sampled points under a certain error criterion or error function, regardless of their location, any subsequent interpolation will incur in a certain error versus the true unknown value

$$\text{Err}(M) = \varepsilon_O^M. \quad (\text{A.1})$$

On the other end, using a minimum of  $k + 1$  non-co-hyperplanar points for interpolation still assumed to be in the optimal position, will yield the largest error for the oracle method;

$$\text{Err}(k + 1) = \varepsilon_O^{k+1}. \quad (\text{A.2})$$

Now let us assume the error function is bounded. While theoretically we can only assume such an upper bound, in practice (regardless of the domain of application) this is utterly circumstantial as no physical sensor has infinite dynamic range and hence, the error between interpolating a value and exactly inverting every point is guaranteed to be bounded in practical terms.

The rationale (but not algebra) for escaping the combinatorial explosion is analogous to a regular limit; being able to come as 'arbitrarily' close to the optimal as needed yet doing it with fewer points. The regular methods (grid-like, gradient) use all  $M$  points yet the loci of such sampling dimensions may not be optimal. This gives an error of  $\varepsilon_B^M \geq \varepsilon_O^M$ . Notwithstanding, by breaking some assumptions undersampling is viable.

We can depart from a C&C case using  $M$  points with a suboptimal error  $\varepsilon_C^M \geq \varepsilon_O^M$ . This is the best a single run of C&C with  $M$  points can perform. Since the error function is bounded and  $\varepsilon_C^M$  is the best C&C can do, then removing one point from such an optimal configuration of C&C gives  $M - 1$  points to be sampled. This will give a new error  $\varepsilon_C^{M-1} \geq \varepsilon_C^M$ . It is trivial from here to construct a sequence of growing error as we remove points  $\varepsilon_C^{k+1} \geq \dots \geq \varepsilon_C^{M-1} \geq \varepsilon_C^M$ . 'Reading' this sequence backwards gives us our escape from the combinatorial explosion; we can use a lower number of points yet still keep the bounded error at bay.

## ORCID iDs

R J Ward  <https://orcid.org/0000-0003-3814-5910>

F Orihuela-Espina  <https://orcid.org/0000-0001-8963-7283>

## References

- Abdalsalam O, Howard S and O'Sullivan T D 2024 Phase-based structured interrogation frequency-domain near-infrared spectroscopy *J. Opt. Soc. Am. A* **41** 1500
- Angelo J, Vargas C R, Lee B T, Bigio I J and Gioux S 2016 Ultrafast optical property map generation using lookup tables *J. Biomed. Opt.* **21** 110501
- Arifler D, Zhu T, Madaan S and Tachtsidis I 2015 Optimal wavelength combinations for near-infrared spectroscopic monitoring of changes in brain tissue hemoglobin and cytochrome *c* oxidase concentrations *Biomed. Opt. Express* **6** 933
- Arridge S R and Hebden J C 1997 Optical imaging in medicine: II. Modelling and reconstruction *Phys. Med. Biol.* **42** 841–53
- Ávila-Sansores S-M, Rodríguez-Gómez G, Tachtsidis I and Orihuela-Espina F 2020 Interpolated functional manifold for functional near-infrared spectroscopy analysis at group level *Neurophotonics* **7** 1–17
- Bellhouse D R 1984 A review of optimal designs in survey sampling *Can. J. Stat.* **12** 53–65
- Dehghani H, Eames M E, Yalavarthy P K, Davis S C, Srinivasan S, Carpenter C M, Pogue B W and Paulsen K D 2009 Near infrared optical tomography using NIRFAST: algorithm for numerical model and image reconstruction *Commun. Numer. Methods Eng.* **25** 711–32
- Delpy D T, Cope M, Zee P v d, Arridge S, Wray S and Wyatt J 1988 Estimation of optical pathlength through tissue from direct time of flight measurement *Phys. Med. Biol.* **33** 1433–42
- Eastmond C, Subedi A, De S and Intes X 2022 Deep learning in fNIRS: a review *Neurophotonics* **9** 041411
- Fantini S, Franceschini M A and Gratton E 1994 Semi-infinite-geometry boundary problem for light migration in highly scattering media: a frequency-domain study in the diffusion approximation *J. Opt. Soc. Am. B* **11** 2128
- Gómez C A and Roblyer D 2024 Effects of skin tone and adipose thickness on frequency domain near-infrared spectroscopy and diffuse correlation spectroscopy *Biophoton. Discovery* **2** 012503
- Hemingway P 2002 n-Simplex interpolation *Technical Report*
- Herrera-Vega J, Treviño-Palacios C G and Orihuela-Espina F 2017 Neuroimaging with functional near infrared spectroscopy: from formation to interpretation *Infrared Phys. Technol.* **85** 225–37
- Jacques S L 2013 Corrigendum: optical properties of biological tissues: a review *Phys. Med. Biol.* **58** 5007–8
- Kao T-C and Sung K-B 2022 Quantifying tissue optical properties of human heads *in vivo* using continuous-wave near-infrared spectroscopy and subject-specific three-dimensional Monte Carlo models *J. Biomed. Opt.* **27** 083021
- Li J and Heap A D 2014 Spatial interpolation methods applied in the environmental sciences: a review *Environ. Modelling Softw.* **53** 173–89
- McBratney A, Webster R and Burgess T 1981 The design of optimal sampling schemes for local estimation and mapping of regionalized variables-I *Comput. Geosci.* **7** 331–4
- Mersereau R 1979 The processing of hexagonally sampled two-dimensional signals *Proc. IEEE* **67** 930–49
- Olea R A 1984 Sampling design optimization for spatial functions *J. Int. Assoc. Math. Geol.* **16** 369–92
- Orihuela-Espina F 2005 Modelling and verification of the diffuse reflectance of the ocular fundus *PhD Thesis* The University of Birmingham Birmingham
- Stillwell R A, Kitsmiller V J, Wei A Y, Chong A, Senn L and O'Sullivan T D 2021 A scalable, multi-wavelength, broad bandwidth frequency-domain near-infrared spectroscopy platform for real-time quantitative tissue optical imaging *Biomed. Opt. Express* **12** 7261
- Wang L, Jacques S L and Zheng L 1995 MCML-Monte Carlo modeling of light transport in multi-layered tissues *Comput. Methods Programs Biomed.* **47** 131–46
- Ward R J, Bale G, Dehghani H and Orihuela-Espina F 2023 Low-cost bNIRS reconstruction of oxygenation and cytochrome-c-oxidase using off the shelf Spectrometers *FNIRS UK 2023* Essex
- Wei X, Wu Y-Z and Chen L-P 2012 A new sequential optimal sampling method for radial basis functions *Appl. Math. Comput.* **218** 9635–46
- Yang B 2021 Rapid quantification of tissue perfusion properties with a two-stage look-up table: a simulation study (available at: <http://biorxiv.org/lookup/doi/10.1101/2021.11.04.467306>)
- Yoo J *et al* 2020 Deep learning diffuse optical tomography *IEEE Trans. Med. Imaging* **39** 877–87
- Zhang N, Canini K, Silva S and Gupta M 2021 Fast linear interpolation *ACM J. Emerg. Technol. Comput. Syst.* **17** 1–15